
CERN

T-systems storage: what?

- 500TB of block storage
- Expect to deploy two systems
  - One shared Atlas/CMS
    - ~35 disk servers
      - ~2x5TB on each
      - High speed “disk”, 20k iop/s, 200MB/s
  - One dedicated to Alice
    - ~15 disk servers

T-systems storage – DPM

- DPM architecture involves
  - Head node
    - Metatdata, auth
  - N disk servers
    - Clients redirected here for i/o
- Needs to be globally accessible
- Currently have a 5 node test system installed
  - Basic tests are passing
    - Some rough edges due to NAT and reverse DNS

Network defines storage

- Data import and access are critical
- Dimensions of the system largely dictated by network characteristics of the data servers
- Network
  - Between VMs – 1Gb/s
    - Means we need lots of disk servers
  - Internet Inbound – 500Mb/s
    - What limits this?
  - Internet Outbound – 300Mb/s
- In all above cases, more network would mean fewer service nodes
  - -> more CPU workers

NAT/DNS

- NAT
  - Our storage systems don’t expect this!
  - Different behaviour for internal/external clients
  - Requires cloud-local DNS
- DNS reverse lookups
  - Our clients need this (unfortunately)
  - For both public and private IPs

Conclusions

- The VMs and block devices work to spec and seem reliable
  - Very good so far!
- There are some architectural constraints which make life a bit harder
- Hope to deploy the full system this week and start to stress it